arXiv:1609.06870v2 [cs.CV] 30 Sep 2016

Distributed Training of Deep Neural Networks: Theoretical and Practical Limits of Parallel Scalability.

Janis Keuper and Franz-Josef Preundt,
Fraunhofer ITWM,
janis.keuper@itwm.fhg.de

ABSTRACT

This paper presents a theoretical analysis and practical evaluation of the main bottlenecks towards a scalable distributed solution for the training of Deep Neural Networks (DNNs). The presented results show, that the current state of the art approach, using data-parallelized Stochastic Gradient Descent (SGD), is quickly turning into a vastly communication bound problem. In addition, we present simple but fixed theoretic constraints, preventing effective scaling of DNN training beyond only a few dozen nodes. This leads to poor scalability of DNN training in most practical scenarios.

1. INTRODUCTION

The tremendous success of Deep Neural Networks (DNNs) [18, 14] in a wide range of practically relevant applications has triggered a race to build larger and larger DNNs [20], which need to be trained with more and more data, to solve learning problems in fast extending fields of applications. However, training DNNs is a compute and data intensive task: current models take several ExaFLOP to compute, while processing hundreds of petabyte of data [20]. Table 1 gives an impression of the compute complexity and shows, that even the latest compute hardware will take days to train the medium sized benchmark networks used in our experiments. While a parallelization of the training problem over up to 8 GPUs hosted in a single compute node can be considered to be the current state of the art, available distributed approaches [4, 15, 1, 2, 7] yield disappointing results [19] in terms of scalability and efficiency. Figure 1 shows representative experimental evaluations, where strong scaling is stalling after only a few dozen nodes. In this paper, we investigate the theoretical and practical constraints preventing better scalability, namely model distribution overheads (in section 2), data-parallelized matrix multiplication (section 3) and training data distribution (section 4).

	CPU	K80	TitanX	KNL
AlexNet:				
time per iteration	2s	0.9s	0.2s [10]	0.6s
time till convergence	250h	112h	25h [10]	75h
GoogLeNet:				
time per iteration	1.3s	0.36s	-	0.32s
time till convergence	361h	100h	-	89h

Table 1: Approximate computation times for AlexNet with batch size $B = 256$ and 450k iterations and GoogLeNet with $B = 32$ and 1000k iterations. KNL (Xeon Phi "Knights Landing") results with MKL17. TitanX with Pascal GPU. See section 1.2.3.

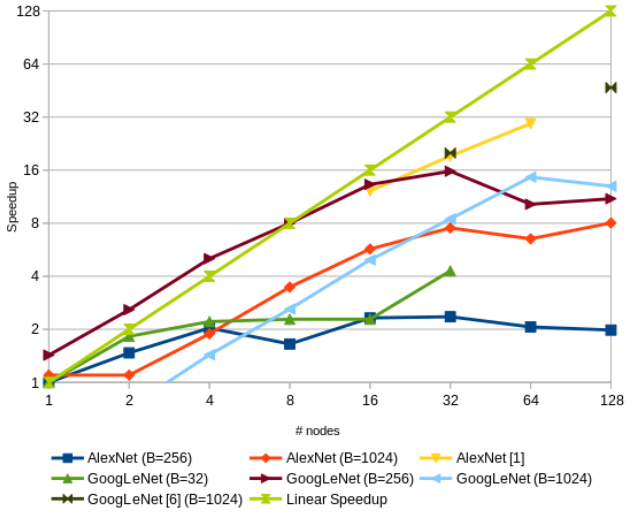

Figure 1: Experimental evaluation of DNN training scalability (strong scaling) for different DNNs with varying global batch sizes B. Results from an "out of the box" installation of *IntelCaffe* on a common HPC system (Details are given in section 1.2)

1.1 Stochastic Gradient Descent

Deep Neural Networks are trained using the *Backpropagation Algorithm* [16]. Numerically, this is formulated as a highly non-convex optimization problem in a very high dimensional space, which is typically solved via Stochastic Gradient Descent (SGD[1]) [3]. SGD, using moderate mini-

[1]Usually, SGD with additional 2nd order terms (moments) are used, but this has no impact on the parallelization.

MLHPC '16

ACM ISBN 123-4567-24-567/08/06.
DOI: 10.475/123_4

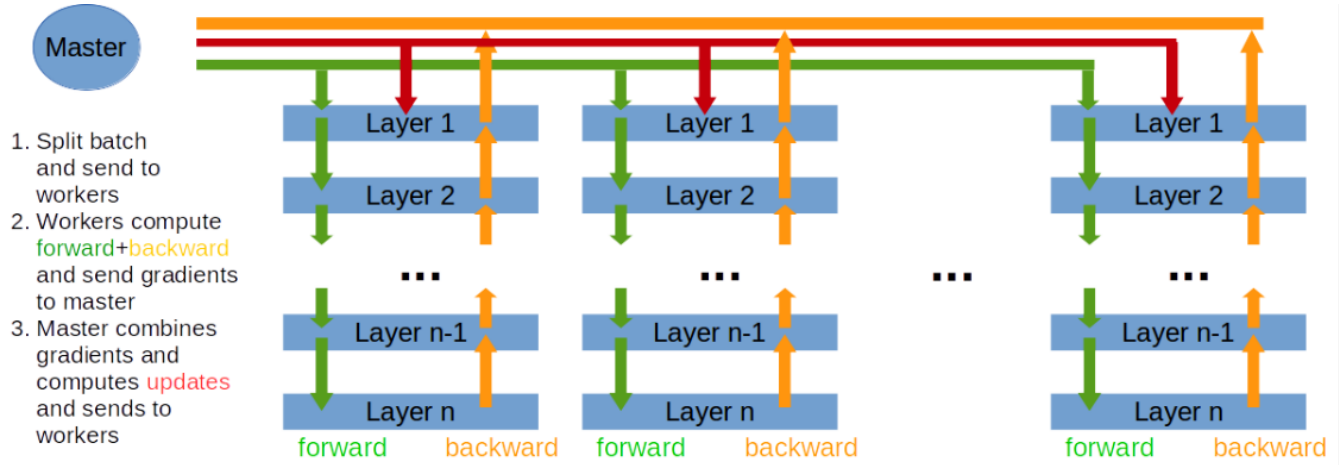

Figure 2: Schematic overview of a distributed SGD implementation of the *Backpropagation Algorithm*.

batch sizes B, provides stable convergence at fair computational costs on a single node. However, it is very hard to parallelize. This is due to the inherently sequential nature of the algorithm, shown in equation 1 and algorithm 1.

$$w_{t+1} \leftarrow w_t - \epsilon \partial_w x_j(w_t), \tag{1}$$

where w_t represents the current sate (e.g. the weights at the neurons), ϵ defines the step-size and $\partial_w x_j(w_t)$ is computed from a given loss-function over the forward results of a small set of training samples (called *mini-batch* and the given training labels). In fact, there are only two ways to speedup SGD: (I) computing updates $\mathbf{\Delta}w$ faster and (II) making larger update steps ϵ. While (I) is hard to achieve in a distributed setting given already low compute times ($<$ 1s) per iteration, (II) is bound by the difficult topologies of the non-convex problems, causing SGD to diverge easily.

1.1.1 Parallelizing SGD

Figure 2 shows the data-parallel version of SGD [4], which is commonly used for single node multi-GPU and distributed implementations: The global batch of B training samples for the current iteration is split into n equal sized sets of size b (with $b = B/n$) of training samples which are then fed to n workers holding synchronous local copies of the model state. The results (gradients) off all workers are then accumulated and used to update the model. Hence, the entire approach is implementing a simple *map-reduce* scheme.
Notably, this scheme implies two different levels parallelization: the data- and task-parallel [12] *Inner Parallelization*, located at the compute units of the nodes using parallel algorithms to compute the forward and backward operations within the layers of the DNN (see section 3 for details on the local parallelization of layer operations), and the *Outer Parallelization* over the distributed batches.

1.2 Experimental Setup

1.2.1 Benchmarks

We apply two widely used convolutional networks (CNNs), *AlexNet* [13] and *GoogLeNet* [22], for the benchmarking of our experimental evaluations. Both neural networks follow different strategies to learn predictive models for the *ImageNet* [17] visual recognition challenge. While *AlexNet* implements a rather shallow network with 3 dominant fully-connected (FC) layers, is *GoogLeNet* using a very deep network with many convolutional layers. Table 2 shows the technical details of both networks.

Algorithm 1 Mini-Batch SGD with samples $X = \{x_0, \ldots, x_m\}$, iterations T, step-size ϵ, batch size B

Require: $\epsilon > 0$
1: **for all** $t = 0 \ldots T$ **do**
2: **randomly draw** batch $M \leftarrow B$ samples from X
3: **Init** $\mathbf{\Delta}w_t = 0$
4: **for all** $x \in M$ **do**
5: **aggregate update** $\mathbf{\Delta}w \leftarrow \partial_w x_j(w_t)$
6: **update** $w_{t+1} \leftarrow w_t - \epsilon \mathbf{\Delta}w_t$
7: **return** w_T

	AlexNet	GoogLeNet
ExaFLOP to convergence	$\sim$ 0.8	$\sim$1.1
# Iterations till convergence	450k	1000k
Model size @32 bit FP	$\sim$250 MB	$\sim$50 MB
Default batch size	256	32
Default step-size	0.01	0.01
# Layers	25	159
# Convolutional layers	5	59
# Fully-connected (FC) layers	3	1
# Weights in FC layers	$\sim$55M	$\sim$1M

Table 2: Properties of the Deep Neural Networks used for the following benchmarks.

1.2.2 Software Framework

We use the MPI based distributed Version (*IntelCaffe*) [1] of the popular open source framework *Caffe* [9] for our evaluation. *IntelCaffe* was built with *CUDA 7.5* and *cuDNN 5* using the latest *Intel* compiler, *MKL*[2] and *IntelMPI*.

[2]Some CPU experiments used the latest DNN extensions of the *MKL17* library which provides special purpose functions for the fast implementation of several layer types like cuDNN for CUDA.

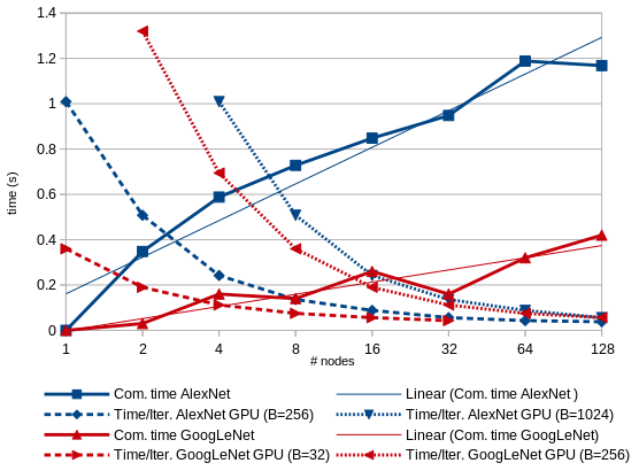

Figure 3: Communication overhead for different models and batch sizes. The scalability stalls when the compute times drop below the communication times, leaving compute units idle. Hence becoming an communication bound problem. Results were generated using a binary tree communication scheme [7].

1.2.3 Hardware

All distributed experiments were conducted on a HPC cluster with nodes holding a dual Xeon E5-2680 v3 CPU (12 cores @ 2.50GHz), a NVIDIA Tesla K80 GPU and FDR-Infiniband interconnects.

2. DISTRIBUTION OVERHEAD

The parallelization of DNN training via SGD (as shown in algorithm 1 and figure 2) requires the communication of the model w_t and the computed gradients $\mathbf{\Delta}w_t$ between all nodes in every iteration t. Since w has to be synchronous in all nodes and $\mathbf{\Delta}w_{t+1}$ can not be computed before w_t is available, the entire communication has to be completed before the next iteration $t+1$. Naturally, one would try to overlap this communication (which can be done layer by layer) with the compute times. However, there are several pitfalls to this strategy: (I) w and $\mathbf{\Delta}w$ have the size of all weights in the neural network, which can be hundreds of megabyte[3], (II) the compute times per iteration[4] are rather low and are decreasing further when scaling to more nodes (see section 3), (III) communication can not start before the forward pass of the network has succeeded (practically cutting the overlay times by half). Ironically, faster compute units (e.g. newer GPUs) in the compute nodes increase the fundamental problem, that the communication time exceeds the compute time, after scaling to only a few nodes. Leaving valuable compute units idle. Figure 3 shows the strong divergence of communication- and compute times. Depending on the model size, the training problem becomes **communication bound** after scaling to only 4 to 8 nodes. This directly correlates to the general scaling results shown in figure 1. Figure 3 also shows, that the network layout has a large impact on the crucial communication/compute ratio: shallow networks with many neurons per layer (like *AlexNet*) scale worse than deep networks with less neurons (like *GoogLeNet*) where longer compute times meet smaller model sizes.

[3]See table 2 for details.
[4]See table 1.

2.1 Limited Network Bandwidth

Limited network bandwidth is one of the key bottlenecks towards the scalability of distributed DNN training. Recently, there have been several approaches proposed to overcome this problem: e.g. [7] introduced a binary communication tree, which reduces the network load to a maximum of $log_2(n)$ peer to peer model/gradient sends at a time. However, expecting linear speedups at the compute side, figure 3 shows that this approach will only move the intersection of the communication/compute ratio by a small factor, as the additional overhead is increasing with the depth of the communication tree.
Other methods try to reduce the model size before the communication. This can be done by (I) a redesign of the network [8] - eliminating unused weights, (II) limiting the numerical precision of the model weights ([6] has shown that one byte per weight is enough), (III) compression (which is available in [1]) or (IV) transmitting only sparse gradient and model information [21]. All these methods have practical impact, moving the scalability by the factor of the model reduction rate. But non of these approaches is able to solve the problem in principle. As model sizes are increasing much faster than the available network bandwidth, the communication overhead remains an unsolved problem.

3. COMPUTATIONAL COSTS AND SCALING OF MATRIX MULTIPLICATIONS

The previously discussed communication overhead is actually a well known problem that has recently been drawing more and more attention in the deep learning community [7, 8, 6, 21]. But communication overhead is not the only problem preventing DNN scalability: there is an even more severe limitation, which turns out to be a hard theoretical constraint. We illustrate this problem by means of a simple

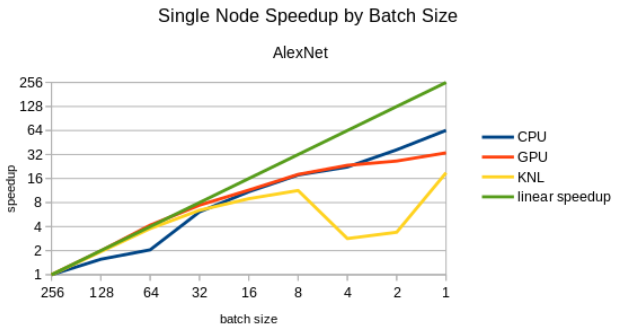

Figure 4: Evaluation of the scalability assuming free communication (simulated by measuring the compute times at a single node at decreasing batch sizes). Results for different compute units.

experiment: assuming that the communication in the distributed DNN training was free, one would expect close to

linear strong scaling[5] properties. However, figure 4 shows, that this is not the case. Again, scalability stalls after only a few nodes. While it is obvious that it is not possible to split the global batch into local batches $b < 1$, thus imposing a **strict scalability limit** at $n = B$, the limitation induced by the batch size are taking effect even for $b >> 1$. To allow further investigation of these results, we provide a layer by layer analysis on the computational complexity and scalability of out benchmark networks.

3.1 Layer by Layer Analysis

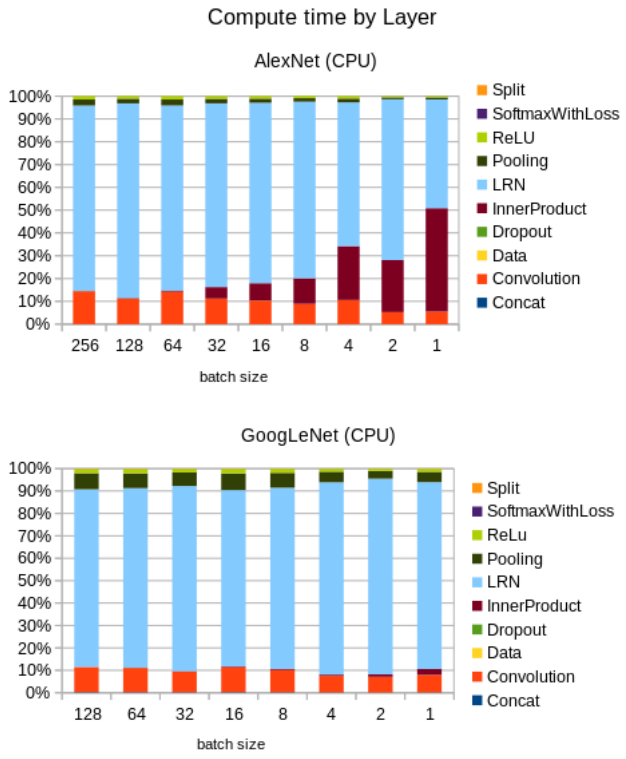

Figure 5: Evaluation of the relative compute time for each layer type (several layers of the same type are accumulated) per training iteration on a single node CPU based system. Top: results for *AlexNet*. Bottom: results for *GoogLeNet*.

Figure 5 shows the analysis for DNN training on CPUs. The dominance of the Local Response Normalization layer (LRN) is caused by a rather poor multi-threaded implementation[6] in *Caffe* and is neglectable in terms of scalability (as shown in figure 5). More interesting is the growing portion of compute time spend in the *InnerProduct* (= Fully-Connected)[7] layer. Figure 6 shows the same tendencies for layer computations on GPUs, where the LRN layer has no significant impact. Yet another interesting observation can be made in figure 15, which shows the impact of the convolution optimization of the cuDNN[8] library used in figure 6. Even more evident than the relative compute portions of the different layer types shown in figures 5, 6, 15 and 14, are the scaling properties by the different layer types. Figure 5 depicts these for DNN training on CPUs (see figure 14 for results on the new Xeon-Phi): all but one layer types show almost perfect linear scaling. Only the significantly compute intensive *InnerProduct* layer scales poorly for batch sizes $b < 64$, which is equivalent to scaling to only $n > 4$ nodes for the original batch size $B = 256$. On the GPU, the crucial *InnerProduct* layer scales much better

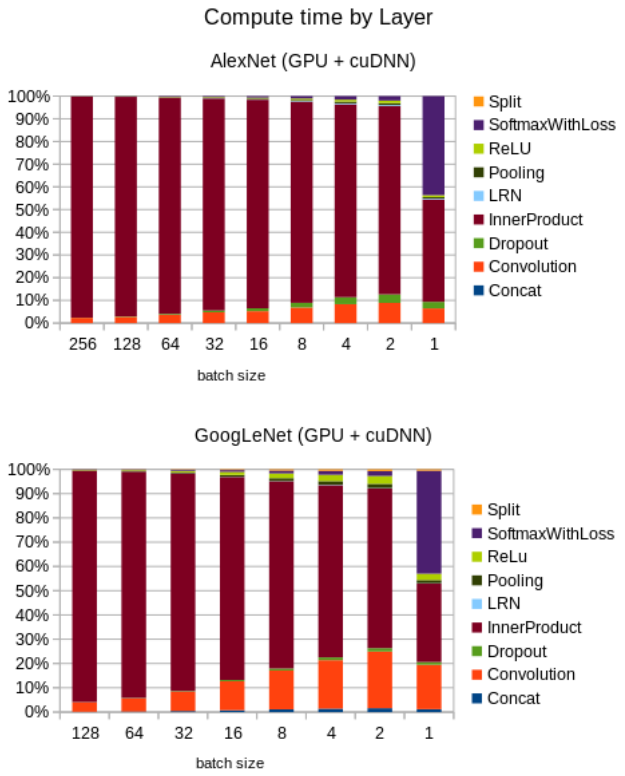

Figure 6: Evaluation of the relative compute time for each layer type (several layers of the same type are accumulated) per training iteration on a single node GPU based system (one K80). Top: results for *AlexNet*. Bottom: results for *GoogLeNet*.

than on the CPU, but still fails linear speedup as we can see accelerations factors around 32x at 256 nodes. Again, the speedup stalls for batch sizes $b \leq 32$.

3.2 Scaling Fully-Connected Layers

The layer by layer analysis revealed the impact of the *Fully-Connected* (FC) layers on the overall scalability. FC layers are the "conventional" neural layers in the deep network architecture, where the actual decision boundaries of a classification problem are modeled. Typically, FC layers hold a large number of neurons which are connected to all inputs. Computationally, FC layers perform a single[9] matrix

[5]because distributed SGD is data-parallel
[6]This problem has been fixed by the MKL17 implementation as shown figure 14.
[7]The changing name convention is due to the naming used in Caffe.
[8]The optimization strategy is also available in MKL17, as shown in figure 14.
[9]Actually there is a second small matrix multiplication for

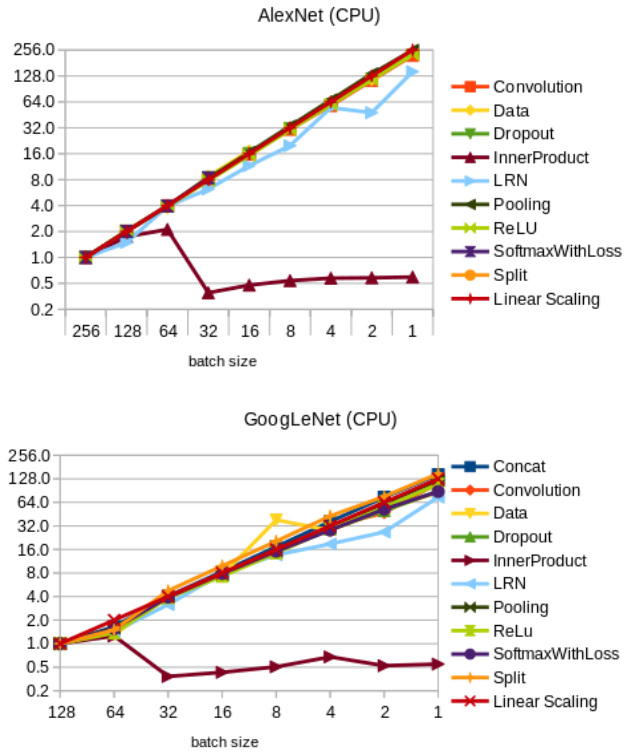

Figure 7: Speedup achieved by reducing the batch size - computed from the results in Fig. 5. Top: results for *AlexNet*. Bottom: results for *GoogLeNet*.

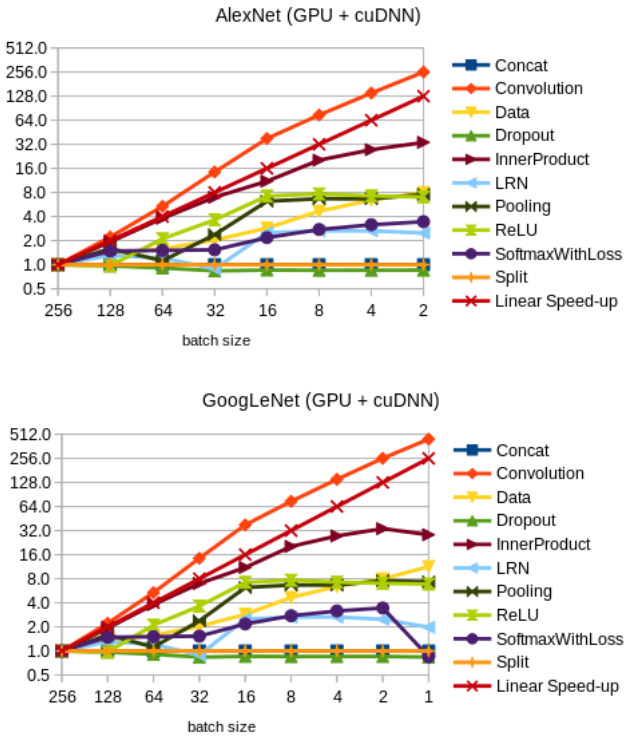

Figure 8: Speedup achieved by reducing the batch size - computed from the results in Fig. 6. Top: results for *AlexNet*. Bottom: results for *GoogLeNet*.

Layer	# operations	matrix sizes
Fully Connected	1	$b \times I * I \times O$
Convolutional	b	$C \times I * I \times Z$
Softmax	b	$I \times 1 * 1 \times 1$

Definitions:

I: Input size from top layer
O: Output size of this layer
b: local Batch size (train or validation)
C: Number of filters
c: Number of input channels (RBG image: $c = 3$)
P: Patch size (i.e. pixel)
k: kernel size
Z: Effective size after kernel application.
For convolution $Z := \left(\sqrt{P} - \lfloor (k/2) \right)^2$

Table 3: Size and number of of the matrix multiplications (sgemm) per forward pass for selected layers.

multiplication per pass. Table 3 shows the impact of the batch size b on the size, shape and number of matrix multiplications. While b only affects the number of matrix operations for *Convolutional* layers (which can implemented task-parallel [12]), it directly reshapes the left-hand matrix in the FC sgemm operation in a very unfavorable way. For typically large I and O (e.g. for a layer in *AlexNet* we find $I = 4096, O = 9192$), $b = B/n$ decreases from $B = 256$ - producing "degenerated" (maximal non-square) matrices. This "degeneration" hurts the *Inner Parallelization* of the matrix multiplication (see section 1.1.1), where the sgemm is either multi-threaded by the MKL Blas Library or parallelized via cudaBlas. Both implementations have an optimal performance for square matrices and suffer from the "degeneration" [5]. Hence, speedups gained by the *Outer Parallelization* (the data-parallel SGD) start to harm the performance of *Inner Parallelizations* and cause a "scalability deadlock". Figure 9 shows the impact of b on the MKL sgemm speedup performance. It is not surprising, that this evaluation shows exactly the same speedup characteristic as the overall communication free scaling experiment in fig. 4.

3.3 Increasing the global Batch Size

A simple way to overcome the stalling scalability beyond 8 nodes (or $b < 32$), has recently been suggested in [7] and is also utilized in [1]: increasing the global batch size B to the extend, that the worker batch size keeps an effective size $b \geq 32$ for the *Inner Parallelization*.
Figure 11 shows, that this strategy is actually providing almost perfect linear speedup up to 128 nodes for a global batch size of 4096. However, these results have to be taken with strong caution: Increasing the global batch size also increases the computational complexity of the problem linearly. Beyond a certain batch size, SGD will not converge significantly faster in terms of the number of iterations. Hence, large batch sizes will increase the computation time per iteration while the number of iterations stays constant.

the computation of the bias which we neglect here.

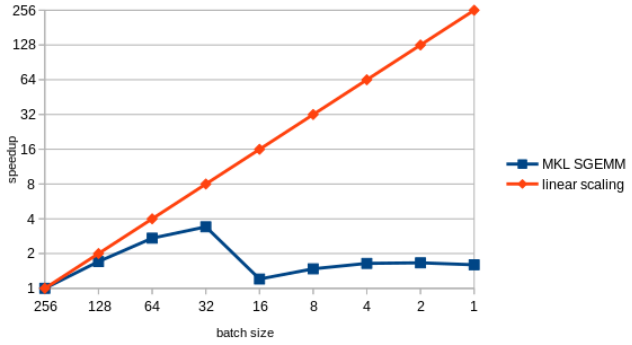

Figure 9: MKL SGEMM: impact of the batch size b on the MKL sgemm speedup performance for matrix multiplications with the shape $b \times 4096 * 4096 \times 9192$. These matrix shapes correspond to the sgemms computed in the largest *Fully-Connected* layer of *AlexNet*.

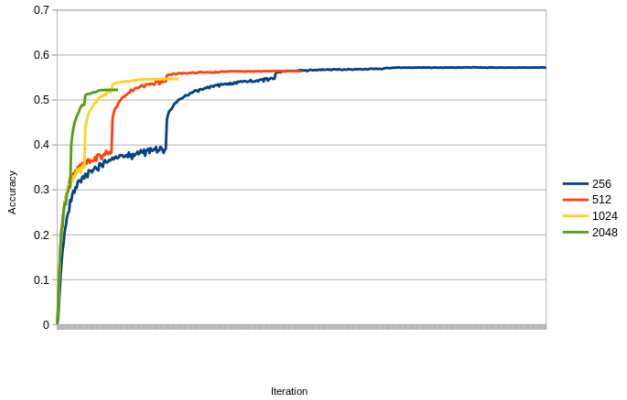

Figure 10: Full validation accuracy plot for *AlexNet* with different large batch sizes. Settings $[B = 256, \epsilon = 0.01, iter = 450k]$, $[B = 512, \epsilon = 0.02, iter = 225k]$, $[B = 1024, \epsilon = 0.04, iter = 112k]$, $[B = 2048, \epsilon = 0.08, iter = 56k]$

In order to reduce the number of iterations till convergence, one would have to increase the step size as well. The authors of [7] argue, that larger batch sizes will provide more stable gradient information which should allow larger step sizes. If it was possible to increase the step size in the same way this is done with the batch size, one would yield perfect linear scaling.

Sadly, this is hardly the case. Figure 10 shows the accuracy plots for *AlexNet*, computed till full convergence with differently large global batch sizes. The experiments were performed on a single KNL node to avoid possible interferences in a distributed setting[10]. The step sizes ϵ were increased according to the batch size as suggested by [7], while the number of iterations has been decreased by the same factor. The results are quite disappointing: while we reach linear speedup as expected, the validation accuracy is suffering

[10]The KNL provides enough memory for such large batch sizes. On common GPUs with 12GB memory, the bach size limit is $b = 256$ for *AlexNet* and $b = 128$ for *GoogLeNet*.

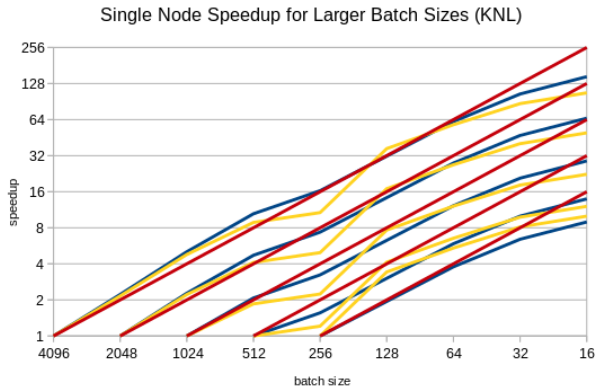

Figure 11: Scaling properties for increased global batch sizes in the "free communication" scenario. Yellow lines show the results for *AlexNet*, blue lines *GoogLeNet* and red lines indicate perfect linear speedup. NOTE: All speedups are computed with respect to the compute time of the enlarged global batch sizes, not the original batch sizes ($B = 256$ for *AlexNet* and $B = 32$ for *GoogLeNet*).

significantly: These experimental results confirm the theoretic analysis by [11], who showed that large batch sizes lead to sharp minima with poorer generalization properties. Considering, that an early stopping of the original problem when reaching the according error rates yields almost the same speedup as the parallelized large batch variants, shows that this approach might not be suitable to solve the scaling problem of the matrix multiplications.

batch size	speedup	step-size	accuracy
256	1	$\epsilon = 0.01$	57.2%
512	2	$\epsilon = 0.02$	56.4%
1024	4	$\epsilon = 0.04$	54.7%
2048	8	$\epsilon = 0.08$	52.2%

Table 4: Effect of choosing larger step-sizes ϵ on the resulting test/validation accuracy. Our experiments show, that larger batch sizes con not compensate for the loss in accuracy.

3.4 Non-Scaling Layers

Figure 8 also shows very poor scalability for some layers like *Dropout, Pooling* or *LRN*. This is mostly due to the fact that these layers are computed so fast, that the latency of loading the data to the GPU becomes the dominant constant factor. Overall, these particular layers consume only a marginal portion of the total compute time[11] (0.1% for *AlexNet* and 0.3% for *GoogLeNet*, assuming that all other layers are parallelizable). Applying *Amdahl's Law* shows in figure 12, that this still affects the scalability in the long run. Again, scalability begins to stall at $n > 32$.

4. PARALLEL TRAINING DATA ACCESS

[11]See figure 6.

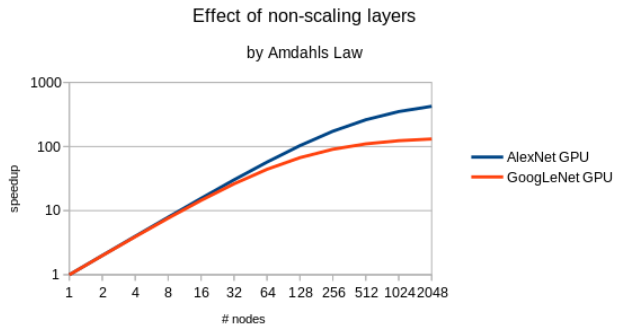

Figure 12: Effect of non-scaling layers to the overall scalability after Amdahl's Law.

So far, the analysis in the previous sections neglected another crucial bottleneck towards scalable distributed DNN training: the distribution of the training data (a.k.a. the batches) to the worker nodes. We specifically avoided this problem in all prior experiments by holding copies of the entire training set on local SSDs on every worker node. However, this approach not only requires the availability of NVRAM (or other high speed local storage) at every node, it is also very inefficient to copy hundreds of gigabyte[12] to each worker node before the actual training can be started.

4.1 Network Bandwidth Revised

Using a centralized storage for the training data, like a database server or a distributed file system, would offer a more convenient and resource effective solution. Compared to the petabytes of communication caused by the distributed SGD (see section 2), the distribution load of training data appears to be neglectable: e.g. for *AlexNet* we have 100 epochs (=full pass of the training data) till convergence, resulting in $100 \times 150\text{GB} = 15\text{TB}$ of total data traffic compared to $450000 \times 250\text{MB} \times 2(n-1)$ in gradient and update communication[13]. But this assumption holds only as long as the SGD communication leaves some bandwidth for the data transfers. Figure 13 shows the practical consequences of storing the training data on a *Lustre* distributed file system[14] when the network bandwidth is exceeded by the SGD communication.

4.2 Small Files - High Speed Random Access

Bandwidth is not the only problem when it comes to the usage of parallel file systems[15]. There are also latency issues, which are caused by the structure of the training data used in many deep learning applications: typically learning samples come as large collections of small files (e.g. images, audio sequences or texts) which need to be accessed at random during the DNN training. Many workers simultaneously polling the file system metadata servers for large numbers of random files easily causes large response times or even the break down of the distributed file system.

[12] *ImageNet 1000* [17] has $\sim$ 150GB of training data, larger real world problems easily exceed many terabytes.
[13] Assuming a parameter server with n workers and 450k iterations.
[14] Storage on SSDs, Interconnect FDR Infiniband.
[15] Which are the standard storage solution on HPC clusters.

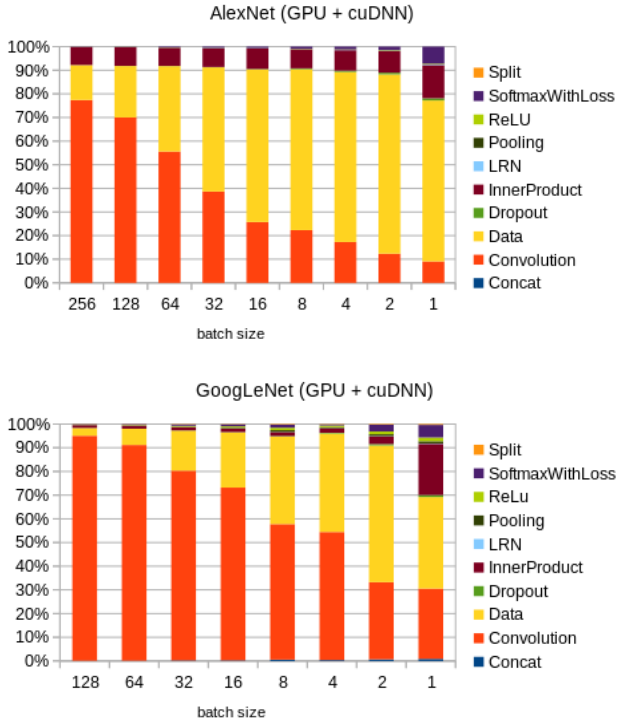

Figure 13: Influence of the Data layer on compute times. This figure shows, that storing the training data on a distributed file system is prune to cause huge performance bottlenecks. Compute units idle during the time spend in the data layer (Compare these results with the compute times shown in Fig. 6 were the training data was stored locally).

5. CONCLUSIONS

In this paper, we analysed and discussed the major theoretical and practical limits of current approaches towards scalable distributed DNN training. We showed three specific bottlenecks, namely the communication overhead, parallelization of matrix operations and training data distribution which need to be solved in order to achieve a sustainably scalable solution which should allow strong scaling to thousands of nodes. Currently, effective scaling is not possible beyond 16 nodes.

6. ACKNOWLEDGMENTS

The authors thank the Center for Information Services and High Performance Computing (ZIH) at TU Dresden for generous allocations of computer time.

7. REFERENCES

[1] https://github.com/intelcaffe/caffe. 2016.
[2] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
[3] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of*

COMPSTAT'2010, pages 177–186. Springer, 2010.

[4] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le, et al. Large scale distributed deep networks. In *Advances in neural information processing systems*, pages 1223–1231, 2012.

[5] J. Demmel, D. Eliahu, A. Fox, S. Kamil, B. Lipshitz, O. Schwartz, and O. Spillinger. Communication-optimal parallel recursive rectangular matrix multiplication. In *Parallel & Distributed Processing (IPDPS), 2013 IEEE 27th International Symposium on*, pages 261–272. IEEE, 2013.

[6] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan. Deep learning with limited numerical precision. *CoRR, abs/1502.02551*, 392, 2015.

[7] F. N. Iandola, K. Ashraf, M. W. Moskewicz, and K. Keutzer. Firecaffe: near-linear acceleration of deep neural network training on compute clusters. *CoRR*, abs/1511.00175, 2015.

[8] F. N. Iandola, M. W. Moskewicz, K. Ashraf, S. Han, W. J. Dally, and K. Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and< 1mb model size. *arXiv preprint arXiv:1602.07360*, 2016.

[9] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.

[10] J. Johnson. https://github.com/jcjohnson/cnn-benchmarks. 2016.

[11] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.

[12] A. Krizhevsky. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997*, 2014.

[13] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[14] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[15] H. Ma, F. Mao, and G. W. Taylor. Theano-mpi: a theano-based distributed training framework. *arXiv preprint arXiv:1605.08325*, 2016.

[16] J. Ngiam, A. Coates, A. Lahiri, B. Prochnow, Q. V. Le, and A. Y. Ng. On optimization methods for deep learning. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 265–272, 2011.

[17] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

[18] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

[19] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu. On parallelizability of stochastic gradient descent for speech dnns. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 235–239. IEEE, 2014.

[20] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[21] R. Spring and A. Shrivastava. Scalable and sustainable deep learning via randomized hashing. *arXiv preprint arXiv:1602.08194*, 2016.

[22] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015.

Appendix

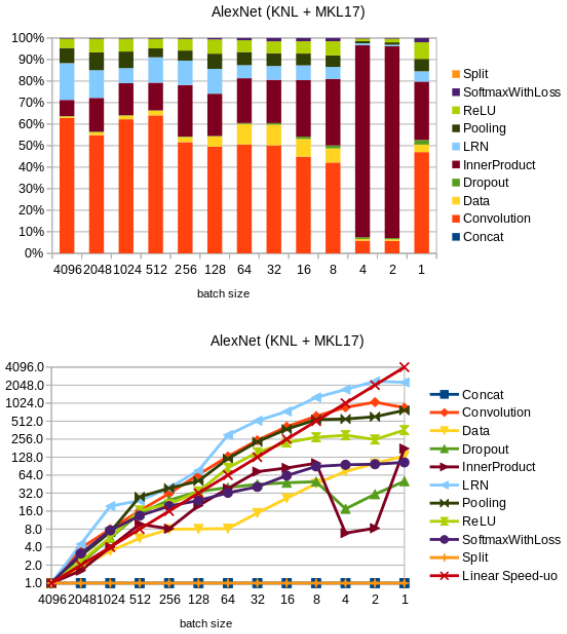

Figure 14: Additional layer by layer results for the KNL (see section 3.1 for details). Top: Proportional compute times by layer type and batch size for *AlexNet*. Bottom: Speedups by layer type ans batch size for *AlexNet*.

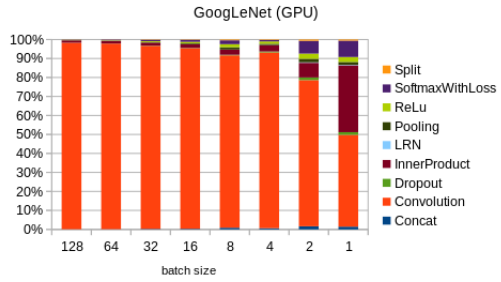

Figure 15: Layer by Layer analysis for *GoogLeNet* without cuDNN.